

WHAT IS CLAIMED IS:

1 1. A method of optimizing instructions included in a program being executed, the
2 method comprising:
3 collecting information describing a frequency of occurrence of a plurality of cache
4 misses caused by at least one instruction;
5 identifying a performance degrading instruction;
6 optimizing the program to provide an optimized sequence of instructions, the
7 optimized sequence of instructions comprising at least one prefetch
8 instruction; and
9 modifying the program being executed to include the optimized sequence.

1 2. The method of claim 1, wherein the program comprises a plurality of sequence
2 of instructions.

1 3. The method of claim 1, wherein the performance degrading instruction
2 contributes to highest frequency of occurrence of the plurality cache misses.

1 4. The method of claim 1, wherein the performance degrading instruction
2 contributes to highest degradation in the program performance.

1 5. The method of claim 1, wherein the at least one instruction is the performance
2 degrading instruction.

1 6. The method of claim 1, wherein optimizing the program comprises inserting
2 the at least one prefetch instruction prior to the performance degrading instruction.

1 7. The method of claim 1, wherein the plurality cache misses are L2/L3 cache
2 misses.

1 8. The method of claim 1, wherein the optimized sequence is prepared while the
2 program is placed in a suspend mode.

1 9. The method of claim 8, wherein modifying the program comprises:
2 changing the program from the suspend mode to the execution mode.

1 10. The method of claim 1, wherein optimizing the program comprises:
2 receiving information describing a dependency graph for the at least one instruction;
3 determining whether a cyclic dependency pattern exists in the dependency graph;
4 if the cyclic dependency pattern exists then, computing stride information derived
5 from the cyclic dependency pattern; and
6 inserting the prefetch instruction derived from the stride information, the prefetch
7 instruction being inserted into the program prior to the performance degrading
8 instruction.

1 11. The method of claim 10, wherein the dependency graph is a backward slice
2 from the performance degrading instruction.

1 12. The method of claim 1, wherein modifying the program comprises:
2 storing the optimized sequence;
3 redirecting a sequence of instructions having the performance degrading instruction to
4 include the optimized sequence.

1 13. A method of optimizing a program comprising a plurality of execution paths,
2 the method comprising:
3 collecting information describing a plurality of occurrences of a plurality of cache
4 miss events during a runtime mode of the program;
5 identifying a performance degrading execution path in the program;
6 modifying the performance degrading execution path to define an optimized
7 execution path, the optimized execution path comprising at least one prefetch
8 instruction;
9 storing the optimized execution path; and
10 redirecting the performance degrading execution path in the program to include the
11 optimized execution path.

1 14. The method of claim 13, wherein the plurality of cache miss events are caused
2 by an execution of a plurality of performance degrading instructions.

1 15. The method of claim 13, wherein identifying the performance degrading path
2 comprises identifying a performance degrading instruction contributing to highest plurality of
3 occurrences of cache miss events.

1 16. The method of claim 13, wherein the optimized execution path is defined
2 while placing the program in a suspend mode from the runtime mode.

1 17. The method of claim 16, wherein the optimized execution path is executed on
2 resuming the runtime mode of the program code from the suspend mode.

1 18. The method of claim 16, wherein redirecting the performance degrading
2 execution path comprises:

3 changing the program mode from the suspend mode to the execution mode.

1 19. The method of claim 13, wherein the performance degrading execution path
2 comprises a performance degrading instruction causing the cache miss event.

1 20. The method of claim 19, wherein the at least one prefetch instruction is
2 inserted prior to the performance degrading instruction.

1 21. The method of claim 13, wherein identifying the performance degrading
2 execution path comprises determining whether a cache miss event of the plurality of cache
3 miss events is an L2/L3 cache miss.

1 22. The method of claim 13, wherein identifying the performance degrading path
2 comprises identifying a performance degrading instruction contributing to highest
3 degradation in the program performance.

1 23. The method of claim 13, wherein modifying the performance degrading
2 execution path comprises:
3 receiving information describing a dependency graph for a program degrading
4 instruction contributing to highest occurrence of the plurality of cache miss
5 events, the performance degrading instruction being included in the
6 performance degrading execution path;
7 determining whether a cyclic dependency pattern exists in the graph;
8 if the cyclic dependency pattern exists then, computing stride information derived
9 from the cyclic dependency pattern; and
10 inserting the at least one prefetch instruction derived from the stride information, the
11 at least one prefetch instruction being inserted into the optimized execution
12 path prior to the performance degrading instruction.

1 24. The method of claim 23, wherein the dependency graph is a backward slice
2 from the performance degrading instruction.

1 25. A method of optimizing a program, the method comprising:
2 receiving information describing a dependency graph for an instruction causing
3 frequent cache misses, the instruction being included in the program;
4 determining whether a cyclic dependency pattern exists in the graph;
5 if the cyclic dependency pattern exists then, computing stride information derived
6 from the cyclic dependency pattern;
7 inserting an at least one prefetch instruction derived from the stride information, the at
8 least one prefetch instruction being inserted into the program prior to the
9 instruction causing the frequent cache misses;
10 reusing the at least one prefetch instruction in the program for reducing subsequent
11 cache misses; and
12 performing said receiving, said determining, said computing, said inserting and said
13 reusing during runtime of the program.

1 26. A computer-readable medium having a computer program accessible
2 therefrom, wherein the computer program comprises instructions for:
3 collecting information describing a frequency of occurrence of a plurality of cache
4 misses caused by at least one instruction;
5 identifying a performance degrading instruction;
6 optimizing the computer program to provide an optimized sequence of instructions,
7 the optimized sequence of instructions comprising at least one prefetch
8 instruction; and
9 modifying the computer program being executed to include the optimized sequence.

1 27. A computer system comprising:
2 a processor;
3 a memory coupled to the processor;
4 a program comprising instructions, the program being stored in memory, the
5 processor executing instructions to:
6 collect information describing a frequency of occurrence of a plurality
7 of cache misses caused by at least one instruction;
8 identify a performance degrading instruction;
9 optimize the program to provide an optimized sequence of instructions,
10 the optimized sequence of instructions comprising at least one
11 prefetch instruction; and
12 modify the program being executed to include the optimized sequence.
13